EMinRET: Heuristic for Energy-Aware VM Placement with Fixed Intervals and Non-preemption

Nguyen Quang-Hung, Nam Thoai

Faculty of Computer Science and Engineering,
HCMC University of Technology, VNU-HCM
Ho Chi Minh City, Vietnam
Email: {hungnq2,nam}@cse.hcmut.edu.vn

arXiv:1511.06825v1 [cs.NI] 21 Nov 2015

***Abstract*—Infrastructure-as-a-Service (IaaS) clouds have become more popular enabling users to run applications under virtual machines. This paper investigates the energy-aware virtual machine (VM) allocation problems in IaaS clouds along characteristics: multiple resources, and fixed interval times and non-preemption of virtual machines. Many previous works proposed to use a minimum number of physical machines; however, this is not necessarily a good solution to minimize total energy consumption in the VM placement with multiple resources, fixed interval times and non-preemption. We observed that minimizing total energy consumption of physical machines is equivalent to minimize the sum of total completion time of all physical machines. Based on the observation, we propose EMinRET algorithm. The EMinRET algorithm swaps an allocating VM with a suitable overlapped VM, which is of the same VM type and is allocated on the same physical machine, to minimize total completion time of all physical machines. The EMinRET uses resource utilization during executing time period of a physical machine as the evaluation metric, and will then choose a host that minimizes the metric to allocate a new VM. In addition, this work studies some heuristics for sorting the list of virtual machines (e.g., sorting by the earliest starting time, or the longest duration time first, etc.) to allocate VM. Using the realistic log-trace in the Feitelson's Parallel Workloads Archive, our simulation results show that the EMinRET algorithm could reduce from 25% to 45% energy consumption compared with power-aware best-fit decreasing (PABFD) [1]) and vector bin-packing norm-based greedy algorithms (VBP-Norm-L1/L2 [2]). Moreover, the EMinRET heuristic has also less total energy consumption than our previous heuristics (e.g. MinDFT and EPOBF) in the simulations (using same virtual machines sorting method).**

***Keywords*-energy efficiency; vm placement; EMinRET; IaaS; total busy time; fixed interval; fixed starting time; heuristic**

I. Introduction

IaaS cloud systems are often built from virtualized data centers [3], [4], [5]. Power consumption in a large-scale data centers requires multiple megawatts [6], [7]. Le et al. [7] estimate the energy cost of a single data center is more than $15M per year. As these data centers scale, they will consume more energy. Therefore, advanced scheduling techniques for reducing energy consumption of these cloud systems are highly concerned for any cloud providers to reduce energy cost. Increasing energy cost and the need to environmental sustainability address energy efficiency is a hot research topic in cloud systems. Energy-aware scheduling of VMs in IaaS cloud is still challenging [8], [7], [9], [10], [11].

Many previous works [4], [1], [2] proved that the virtual machine allocation is NP-Hard and proposed to address the problem of energy-efficient scheduling of VMs in cloud data centers. They [4], [1], [2] present techniques for consolidating virtual machines in cloud data centers by using bin-packing heuristics (such as First-Fit Decreasing [2], and/or Best-Fit Decreasing [1]). They attempt to minimize the number of running physical machines and to turn off as many idle physical machines as possible. Consider a d-dimensional resource allocation where each user requests a set of virtual machines (VMs). Each VM requires multiple resources (such as CPU, memory, and IO) and a fixed quantity of each resource at a certain time interval. Under this scenario, using a minimum of physical machines may not be a good solution. Our observations are that using a minimum number of physical machines is not necessarily a good solution to minimize total energy consumption. In a homogeneous environment where all physical servers are identical, the power consumption of each physical server is linear to its CPU utilization, i.e., a schedule with longer working time will consume more energy than another schedule with shorter working time.

TABLE I
EXAMPLE SHOWING THAT USING A MINIMUM NUMBER OF PHYSICAL SERVERS IS NOT OPTIMAL. (*: NORMALIZED DEMAND RESOURCES TO PHYSICAL SERVERS CAPACITY RESOURCES)

VM ID	CPU*	RAM*	Network*	Starttime	Dur. (hour)
VM1	0.5	0.1	0.2	0	10
VM2	0.5	0.5	0.2	0	2
VM3	0.2	0.4	0.2	0	2
VM4	0.2	0.4	0.2	0	2
VM5	0.1	0.1	0.1	0	2
VM6	0.5	0.5	0.2	1	9

To the best of our knowledge, our work is the first work that studies increasing time and resource efficiency-based approach to allocate VMs onto physical machines in other that it minimizes total energy consumption of all physical machines. Each VM requests resource allocation in a fixed starting time and non-preemption for the duration time. We present here an example to demonstrate our ideas to minimize

total energy consumption of all physical machines in the VM placement with fixed starting time and duration time. For example, given six virtual machines (VMs) with their resource demands described in Table I. Our constraints in the example are the maximum capacity of each resource is 1. In the example, a bin-packing-based algorithm could result in a schedule S_1 in which two physical servers are used: one for allocating VM1, VM3, VM4, and VM5; and another one for allocating VM2 and VM6. The resulted total completion time is (10 + 10) = 20 hours. However, in another schedule S_2 in which where VMs are placed on three physical servers, VM1 and VM6 on the first physical server, VM3, VM4 and VM5 on the second physical server, and VM2 on the third physical server, then the total completion time of the five VMs is only (10 + 2 + 2) = 14 hours.

The EMinRET heuristic places VMs that request multiple resources in the fixed interval time and non-preemption into physical machines to minimize total energy consumption of physical machines while meeting all resource requirements. Using numerical simulations, we compare the EMinRET with the popular modified best-fit decreasing (PABFD) [1], two vector bin-packing norm-based greedy (VBP-Norm-L1/L2) [2], and our previous algorithms (e.g. EPOBF-ST/FT [9], and MinDFT-ST/FT [12]). Using two real log-traces (i.e., [13] and [14]) in the Feitelson's Parallel Workloads Archive, our simulation results show that the EMinRET heuristic with its configurations could reduce from 25% to 45% total energy consumption compared with Power-Aware Best-Fit Decreasing (PABFD) [1] and two norm-based greedy heuristics (VBP-Norm-L1/L2) [2]. Moreover, the EMinRET heuristic has also less total energy consumption than our previous heuristics (e.g. MinDFT and EPOBF) in the simulations (using same virtual machines sorting method).

The rest of this paper is structured as follows. Section II discusses related works. Section III describes the energy-aware VM allocation problem with multiple requested resources, fixed starting and duration time. We also formulate the objective of scheduling, and present our theorems. The proposed EMinRET algorithm present in Section IV. Section V discusses our performance evaluation using simulations. Section VI concludes this paper and introduces future works.

II. Related Works

Many previous research [1], [4], [15], [10], [16] proposed algorithms that consolidate VMs onto a small set of physical machines (PMs) in virtualized datacenters to minimize energy/power consumption of PMs. A group in Microsoft Research [2] has studied first-fit decreasing (FFD) based heuristics for vector bin-packing to minimize number of physical servers in the VM allocation problem. Some other works also proposed meta-heuristic algorithms to minimize the number of physical machines. Beloglazov et al. [1], [4] have proposed VM allocation problem as bin-packing problem and presented a power-aware best-fit decreasing (denoted as PABFD) heuristic. PABFD sorts all VMs in a decreasing order of CPU utilization and tends to allocate a VM to an active physical server that would take the minimum increase of power consumption. Knauth et al. [15] proposed the OptSched scheduling algorithm to reduce cumulative machine up-time (CMU) by 60.1% and 16.7% in comparison to a round-robin and First-fit. The OptSched uses an minimum of active servers to process a given workload. In a heterogeneous physical machines, the OptSched maps a VM to a first available and the most powerful machine that has enough VM's requested resources. Otherwise, the VM is allocated to a new unused machine. In the VM allocation problem, however, minimizing the number of used physical machines is not equal to minimizing total of total energy consumption of all physical machines. Previous works do not consider multiple resources, fixed starting time and non-preemptive duration time of these VMs. Therefore, it is unsuitable for the power-aware VM allocation considered in this paper, i.g. these previous solutions can not result in a minimized total energy consumption for VM placement problem with certain interval time while still fulfilling the quality-of-service.

Chen et al [16] observed there exists VM resource utilization patterns. The authors presented an VM allocation algorithm to consolidate complementary VMs with spatial and temporal-awareness in physical machines. They introduce resource efficiency and use norm-based greedy algorithm, which is similar to in [2], to measure distance of each used resource's utilization and maximum capacity of the resource in a host. Their VM allocation algorithm selects a host that minimizes the value of this distance metric to allocate a new VM. Our proposed EMinRET uses a different metric that unifies both increasing time and resource efficiency. In our proposed metric, the increasing time is the difference between two completion time of a host after and before allocating a VM. In addition, our proposed EMinRET core algorithm has swapping step of overlapped VMs together to minimize total completion times of all physical machines.

Some other research [8], [7], [10] considered HPC applications/jobs in HPC clouds. Garg et al. [8] proposed a meta-scheduling problem to distribute HPC applications to cloud systems with distributed N data centers. The objective of scheduling is minimizing CO_2 emission and maximizing the revenue of cloud providers. Le et al. [7] distribute VMs across distributed cloud virtualized data centers whose electricity prices are different in order to reduce the total electricity cost. Takouna et. al., [10] presented power-aware multi-core scheduling and their VM allocation algorithm selects a host which has the minimum increasing power consumption to assign a new VM. The VM allocation algorithm, however, is similar to the PABFDs [1] except that it concerns memory usage in a period of estimated runtime for estimating the host's energy. The work also presented a method to select optimal operating frequency for a (DVFS-enabled) host and configure the number of virtual cores for VMs. Our proposed EMinRET algorithm that differs from these previous works. Our EMinRET algorithm use the VM's fixed starting time and duration time to minimize the total working time on physical servers, and consequently minimize the total energy consumption in

all physical servers. To the best of our knowledge, no existing works that surveyed in [17], [18], [19], [20] have thoroughly considered these aspects in addressing the problem of VM placement.

In 2007, Kovalyov et al. [21] has presented a work to describe characteristics of a fixed interval scheduling problem in which each job has fixed starting time, fixed processing time, and is only processed in the fixed duration time on a available machine. The scheduling problem can be applied in other domains. Angelelli et al. [22] considered interval scheduling with a resource constraint in parallel identical machines. The authors proved the decision problem is NP-complete if number of constraint resources in each parallel machine is a fixed number greater than two.

III. Problem Description

A. Notations

We use the following notations in this paper:
vm_i: The i^{th} virtual machine to be scheduled.
M_j: The j^{th} physical server.
S: A feasible schedule.
P_j^{idle}: Idle power consumption of the M_j.
P_j^{max}: Maximum power consumption of the M_j.
$P_j(t)$: Power consumption of the (M_j) at a time point t.
ts_i: Fixed starting time of vm_i.
dur_i: Duration time of vm_i.
T: Maximum schedule length, which is the time that the last virtual machine will be finished.
$n_j(t)$: Set of indexes of all virtual machines that are assigned to the M_j at time t.
T_j : Total busy time (working time) of the M_j.
e_i: Energy consumption for running the vm_i in the physical machine that the vm_i is allocated.

B. Power consumption model

In this paper, we use the following energy consumption model proposed in [6] for a physical machine. The power consumption of the M_j, denoted as $P_j(.)$, is formulated as follow:

$$P_j(t) = P_j^{idle} + (P_j^{max} - P_j^{idle})U_j(t) \quad (1)$$

The CPU utilization of the physical server at time t, denoted as $U_j(t)$, is defined as the average percentage of total of allocated computing powers of $n_j(t)$ VMs that is allocated to the M_j. We assume that all cores in CPU are homogeneous, i.e. $\forall c = 1, 2, ..., PE_j : MIPS_{j,c} = MIPS_{j,1}$, The CPU utilization is formulated as follow:

$$U_j(t) = (\frac{1}{PE_j \times MIPS_{j,1}}) \sum_{c=1}^{PE_j} \sum_{i \in n_j(t)} mips_{i,c} \quad (2)$$

The energy consumption of the server in period time $[t_1, t_2]$ is formulated as follow:

$$E_j = \int_{t_1}^{t_2} P_j(U_j(t))dt \quad (3)$$

where:
$U_j(t)$: CPU utilization of the M_j at time t and $0 \leq U_j(t) \leq 1$.
PE_j : Number of processing elements (i.e. cores) of the M_j.
$mips_{i,c}$: Allocated MIPS of the cth processing element to the vm_i by the M_j.
$MIPS_{j,c}$: Maximum capacity computing power (Unit: MIPS) of the cth processing element on the M_j.

C. Problem formulation

Given a set of virtual machines vm_i ($i = 1, 2, ..., n$) to be scheduled on a set of physical servers M_j ($j = 1, 2, ..., m$). Each VM is represented as a d-dimensional vector of demand resources, i.e. $vm_i = (x_{i,1}, x_{i,2}, ..., x_{i,d})$. Similarly, each physical machine is denoted as a d-dimensional vector of capacity resources, i.e. $M_j = (y_{j,1}, y_{j,2}, ..., y_{j,d})$. We consider types of resources such as processing element (core), computing power (Million instruction per seconds -MIPS), physical memory (RAM), network bandwidth (BW), and storage. Each vm_i is started at a fixed starting time (ts_i) and is non-preemptive during its duration time (dur_i).

We assume that the power consumption model is linear to CPU utilization. Even if all physical servers are identical and all VMs are identical too, the scheduling is still NP-hard with $d \geq 1$ [2]. With the problem considered in this paper, all physical servers are identical and their power consumption models are linear to their CPU utilization as can be seen in the two equations (1) and (3). The energy consumption of a physical server in a time unit is denoted as E_0 and is the same for all physical servers since the servers are identical. The objective is to find out a feasible schedule S that minimizes the total energy consumption in the equation (4) with $i \in \{1, 2, ..., n\}$, $j \in \{1, 2, ..., m\}$, $t \in [0; T]$ as following:

$$\textbf{Minimize } (E_0 \times \sum_{j=1}^{m} T_j + \sum_{i=1}^{n} e_i) \quad (4)$$

where the total busy time of a physical server [23], denoted as T_j, is defined as length of union of interval times of all VMs that are allocated to a physical machine M_j at time T.

$$T_j = span(\bigcup_{vm_i \in M_j} [ts_i, ts_i + dur_i]) \quad (5)$$

The union of two time intervals [a,b] and [c,d] is defined as: $[a, b] \cup [c, d] = \{x \in \mathbb{R} \mid x \in [a, b] \ or \ x \in [c, d]\}$ and given a time interval $I = [a, b], a \leq b : span(I) = b - a$. If two time interval are not overlapped then span of the two non-overlapped interval is sum of span of each interval, $I_1 = [a, b], I_2 = [c, d], a \leq b \leq c \leq d : span(I_1 \cup I_2) = span(I_1) + span(I_2) = (b - a) + (d - c)$.

The scheduling problem has the following hard constraints that are described in our previous work [12].

D. Theorems

Theorem 1: Given a cloud system with a set of identical physical machines, assume that power consumption of a

physical machine is $P(u) = b + au$, in which: $b = P_{idle}$ is the idle power consumption, $a = P_{max} - P_{idle}$ is the maximum power consumption, and u is the CPU utilization in percentage ($0 \leq u \leq 1$). We denote e_{ij} is energy consumption of each virtual machine i-th that is scheduled on any physical machine j-th. If the u of the mapped virtual machine is a constant, then the energy consumption of each virtual machine, e_{ij}, is independent of any mapping (i.e. any schedule). We have $\forall i \in 1, 2, ..., n, j \in 1, 2, ..., m : e_{ij} = e_i$.

Proof: Recall that the energy consumption is formulated in Equation (3), and power consumption, $P(u)$, is a linear function of CPU utilization, u. Therefore for all $i \in 1, 2, ..., n$ and $j \in 1, 2, ..., m$, we see that e_{ij} is the integral of the $P(u)$ over any time interval [t_1, t_2], and is the same value, denoted as e_i. ■

From Theorem 1, we can imply the following theorem to prove that Equation 4 is equivalent to Equation 6.

Theorem 2: Minimizing total energy consumption is equivalent to minimizing the sum of total busy time of all physical machines ($\sum_{j=1}^{m} T_j$).

$$\textbf{Minimize } (E_0 \times \sum_{j=1}^{m} T_j + \sum_{i=1}^{n} e_i) \sim \textbf{Minimize } (\sum_{j=1}^{m} T_j) \quad (6)$$

Proof: According to the objective function described in (4), E_0 is constant while e_i is independent of any mapping (i.e. any schedule). ■

Based on the above observations, we propose our energy-aware algorithms denoted as EMinRET which is presented in the next section.

IV. EMinRET: Energy-Aware Minimizing Resource Efficiency - Time Heuristic

A. Scheduling algorithm

In this section, we present our energy-aware scheduling algorithm, namely, EMinRET. EMinRET presents a metric to unify the increasing time and estimated resource efficiency when mapping a VM onto a physical machine. Then, EMinRET will choose a host that has minimizing of the metric. Our previous MinDFT-ST/FT [12] only focused on minimizing the increasing time when mapping a VM onto a physical machine. The EMinRET additionally considers resource efficiency during an execution period of a physical machine in order to fully utilize resources in a physical machine. Furthermore, the core EMinRET algorithm can swap an overlapped VM, which has already been assigned to an active physical machine before, with new VM to minimize total completion time of the physical machine. In this paper, two VMs are overlapped if $ts_1 < ts_2 < (ts_1 + dur_1) < (ts_2 + dur_2)$, where ts_1, ts_2, dur_1, dur_2 are starting times and duration times of two VMs. The core EMinRET algorithm will swap new VM and its overlapped VM together if two VMs meet conditions: (i) both VMs are of the same VM type (i.e. the same amount of requested resources such as number of CPU core, physical memory, network bandwidth, storage, etc.); (ii) the new VM has duration time longer than its overlapped VM.

Algorithm 1 EMinRET: Energy-Aware Minimizing Resource Efficiency - Time

```
1: function EMinRET
2:     Input: vmList - a list of virtual machines to be scheduled, hostList - a list of physical servers
3:     Output: mapping (a feasible schedule) or null
4:     vmList = sortVmListByOrder( vmList, order=[starttime, finishtime] )  ▷ 1
5:     m = hostList.size(); n = vmList.size();
6:     T[j] = 0, ∀j ∈ [1, m]
7:     for i = 1 to n do  ▷ on the VMs list
8:         vm = vmList.get(i)
9:         allocatedHost = null
10:        T1 = sumTotalHostCompletionTime( T )
11:        minRETime = +∞
12:        for j = 1 to m do  ▷ on the hosts list
13:            host = hostList.get(j)
14:            hostVMList = sortVmListByOrder( host.getVms(), order=[starttime, finishtime])
15:            for (Vm vmTemp : hostVMList) do
16:                if canSwap(vm, vmTemp) then
17:                    swap(vm, vmTemp) and break
18:                end if
19:            end for
20:            if host.checkAvailableResource( vm ) then
21:
22:                preTime = T[ host.id ]
23:                T[ host.id ] = host.estimateHostTotalCompletionTime( vm )
24:                T2 = sumTotalHostCompletionTime( T )
25:                RETime = EstimateMetricTimeResEff( T2 - T1, host )
26:                if (minRETime > RETime ) then
27:                    minRETime = RETime
28:                    allocatedHost = host
29:                end if
30:                T[ host.id ] = preTime  ▷ Next iterate over hostList and choose the host that minimize the value of different time and resource efficiency
31:            end if
32:        end for
33:        if (allocatedHost != null) then
34:            allocate the vm to the host
35:            add the pair of vm (key) and host to the mapping
36:        end if
37:    end for
38:    return mapping
39: end function
40: sumTotalHostCompletionTime(T[]) = ∑_{j=1}^{m} T_j  ▷ T[1...m]: Array of total completion times of m physical servers
```

Algorithm 2 Estimating the metric for increasing time and resource efficiency

function ESTIMATEMETRICTIMERESEFF
Input: difftime - a different time, host - a candidate physical machine
Output: ret - a value of metric time and resource efficiency
Set $\mathcal{R}$={cpu, ram, netbw, io, storage, time}
j=host.getId(); n_j=host.getVMList();
for $r \in \mathcal{R}$ **do**
Calculate the resource utilization, $U_{j,r}$ as in the Equaltion (7).
end for
$weights[] \leftarrow$ Read resource weights from configuration file.
$ret = (difftime \times weights[time])^2 + \sum_{r \in \mathcal{R}}((1 - U_{j,r}) \times weights[r])^2$ ▷ weights[time] is weight of the different time
return ret
end function

Our previous MinDFT-ST/FT [12] do not have this swapping step, and neither the EPOBF-ST/FT [9] have these steps.

Based on Equation 2, the utilization of a resource r (resource r can be CPU, physical memory, network bandwidth, storage, etc.) of the M_j, denoted as $U_{j,r}$, is formulated as:

$$U_{j,r} = \sum_{s \in n_j} \frac{V_{s,r}}{H_{j,r}}. \quad (7)$$

where n_j is the list of virtual machines that are assigned to the M_j, $V_{s,r}$ is the amount of requested resource r of the virtual machine s (note that in our study the value of $V_{s,r}$ is fixed for each user request), and $H_{j,r}$ is maximum capacity of resource r in the M_j.

Inspired by the work from Microsoft research team [2], [16], resource efficiency of a physical machine j-th, denoted by RE_j, is Norm-based distant [2] of two vectors: normalized resource utilization vector and unit vector **1**, the resource efficiency is formulated as:

$$RE_j = \sum_{r \in \mathcal{R}}((1 - U_{j,r}) \times w_r)^2 \quad (8)$$

where $\mathcal{R}$={cpu, ram, netbw, io, storage}: set of resource types in a host, w_r is weight of resource r in a physical machine.

In this paper, we propose a unified metric for increasing time and resource efficiency that is calculated as:

$$RET = (t^{diff} \times w_{r=time})^2 + \sum_{r \in \mathcal{R}}((1 - U_{j,r}) \times w_r)^2 \quad (9)$$

EMinRET chooses a physical host that has a minimum value of the RET metric to allocate for a VM. We present the pseudo-code of EMinRET in Algorithm 1. The EMinRET can sort the list of VMs by earliest starting time first, or earliest finishing time first, or longest duration time first, etc.. The EMinRET solves the scheduling problem in time complexity of $\mathcal{O}(n \times m \times q)$ where n is the number of VMs to be scheduled, m is the number of physical machines, and q is the maximum number of allocated VMs in the physical machines $M_j, \forall j = 1, 2, ..., m$.

V. PERFORMANCE EVALUATION

A. Algorithms

In this section, we study the following VM allocation algorithms:

- PABFD, a power-aware and modified best-fit decreasing heuristic [1][4]. The PABFD sorts the list of VM_i (i=1, 2,..., n) by their total requested CPU utilization, and assigns new VM to any host that has a minimum increase in power consumption.
- VBP-Norm-LX, a family of vector packing heuristics that is presented as Norm-based Greedy with degree X=1, 2 [2]. Weights of these Norm-based Greedy heuristics use FFDAvgSum which are exp(x), which is the value of the exponential function at the point x, where x is average of sum of demand resources (e.g. CPU, memory, storage, network bandwidth, etc.). VBP-Norm-LX assigns new VM to any host that has minimum of these norm values.
- EPOBF-ST and EPOBF-FT, presented in [9]. The EPOBF-ST and EPOBF-FT algorithms sorts the list of VM_i (i=1, 2,..., n) by their starting time (ts_i) and respectively by their finished time ($ts_i + dur_i$). Both EPOBF-ST and EPOBF-FT choose a host that has maximum of performance-per-watt to assign a new VM. The performance-per-watt is ratio of total of maximum capacity MIPS and maximum host's power consumption.
- MinDFT-ST and MinDFT-FT, presented in [12]. The MinDFT-ST and MinDFT-FT algorithms sorts the list of VM_i (i=1, 2,..., n) by their starting time (ts_i) and respectively by their finished time ($ts_i + dur_i$). Both MinDFT-ST and MinDFT-FT allocate each VM (in a given set of VMs) to a host that has a minimum increase in total completion times of hosts.
- EMinRET, our proposed algorithm discussed in Section IV-A. We evaluate the EMinRET with some its configurations: The EMinRET-1 sorts the list of VMs by VM's earliest starting time first and host's allocated VMs by its starting time. The EMinRET-2 sorts the list of VMs by VM's earliest starting time first and host's allocated VMs by its finishing times. The VM's finishing time, which is sum of its starting time and its duration time, is calculated by ($ts_i + dur_i$). The EMinRET-3 sorts the list of VMs by VM's earliest finishing time first and host's allocated VMs by its starting time. The EMinRET-4 sorts the list of VMs by VM's earliest finishing time first and host's allocated VMs by its finishing time. The EMinRET-5 sorts the list of VMs by VM's longest duration time first and host's allocated VMs by its starting time. The EMinRET-6 sorts the list of VMs by VM's longest duration time first and host's allocated VMs by its finishing time. The EMinRET-7 sorts the list of VMs

by VM's latest finishing time first and host's allocated VMs by its starting time. The EMinRET-8 sorts the list of VMs by VM's latest finishing time first and host's allocated VMs by its finishing time.

B. Methodology

TABLE II
FOUR VM TYPE IN SIMULATIONS

Type	Cores	MIPS	Mem. (MB)	Net. (Mb/s)	Storage (GB)
VM1	2	2500	871	100	5
VM2	1	2000	3840	100	5
VM3	1	1000	1536	100	5
VM4	1	500	613	100	5

TABLE IV
RESULT OF SIMULATIONS USING THE FIRST 300 JOBS OF THE HPC2N SETH LOG-TRACE [14]

Algorithm	#Hosts	#VMs	Energy (KWh)	Saving (%)
(1) PABFD [1] (baseline)	10000	5687	3325.72	0%
(2) VBP Norm L1 [2]	10000	5687	3328.81	0%
(3) VBP Norm L2 [2]	10000	5687	3328.81	0%
(4) EPOBF-ST [9]	10000	5687	2763.78	17%
(6) EPOBF-FT [9]	10000	5687	2786.90	16%
(7) MinDFT-ST [12]	10000	5687	2651.58	20%
(8) MinDFT-FT [12]	10000	5687	2786.90	16%
(9) EMinRET-1	10000	5687	2491.70	25%
(10) EMinRET-2	10000	5687	2490.81	25%
(11) EMinRET-3	10000	5687	2457.73	26%
(12) EMinRET-4	10000	5687	2457.73	26%
(13) EMinRET-5	10000	5687	1957.67	41%
(14) EMinRET-6	10000	5687	1957.99	41%
(15) EMinRET-7	10000	5687	1817.62	45%
(16) EMinRET-8	10000	5687	1817.62	45%

TABLE V
RESULT OF SIMULATIONS USING THE FIRST 50 JOBS OF SDSC BLUE LOG-TRACE [13].

Algorithm	#Hosts	#VMs	Energy (KWh)	Saving (%)
(1) PABFD [1] (baseline)	5000	8368	729.51	0%
(2) VBP Norm-L1 [2]	5000	8368	784.04	-7%
(3) VBP Norm-L2 [2]	5000	8368	784.04	-7%
(4) EPOBF-ST [9]	5000	8368	576.59	21%
(5) EPOBF-FT [9]	5000	8368	649.23	11%
(6) MinDFT-ST [12]	5000	8368	563.27	23%
(7) MinDFT-FT [12]	5000	8368	649.23	11%
(8) EMinRET-1	5000	8368	447.84	39%
(9) EMinRET-2	5000	8368	447.84	39%
(10) EMinRET-3	5000	8368	447.84	39%
(11) EMinRET-4	5000	8368	447.84	39%
(12) EMinRET-5	5000	8368	447.84	39%
(13) EMinRET-6	5000	8368	447.84	39%
(14) EMinRET-7	5000	8368	447.84	39%
(15) EMinRET-8	5000	8368	447.84	39%

We evaluate these algorithms by simulation using the CloudSim [24] to create a simulated cloud data center system that has identical physical machines, heterogeneous VMs, and with thousands of CloudSim's cloudlets [24] (we assume that each HPC job's task is modeled as a cloudlet that is run on a single VM). The information of VMs (and also cloudlets) in these simulated workloads is extracted from two real log-traces (SDSC Blue Horizon log-trace [13]) and HPC2N Seth log-trace [14]) in Feitelson's Parallel Workloads Archive (PWA) [25] to model HPC jobs. When converting from the log-trace, each cloudlet's length is a product of the system's processing time and CPU rating (we set the CPU rating is equal to included VM's MIPS). We convert job's submission time, job's start time (if the start time is missing, then the start time is equal to sum of job's submission time and job's waiting time), job's request run-time, and job's number of processors in job data from the log-trace in the PWA to VM's submission time, starting time and duration time, and number of VMs (each VM is created in round-robin in the four types of VMs in Table II on the number of VMs). Four types of VMs as presented in the Table II is similar to categories in Amazon EC2's VM instances: high-CPU VM, high-memory VM, small VM, and

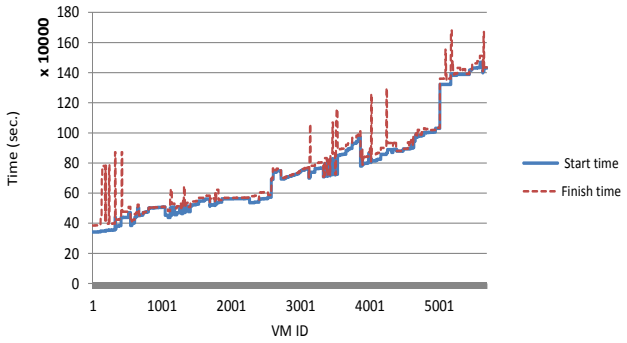

Fig. 1. Starting time (blue line) and finishing time (dotted red line) of VMs in simulations with HPC2N Seth log-trace [14].

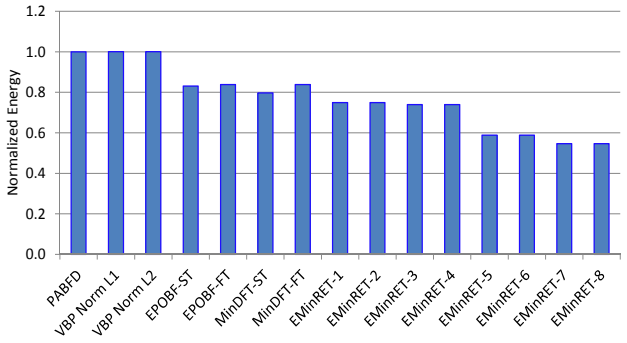

Fig. 2. Normalized energy. Result of simulations with HPC2N Seth log-trace.

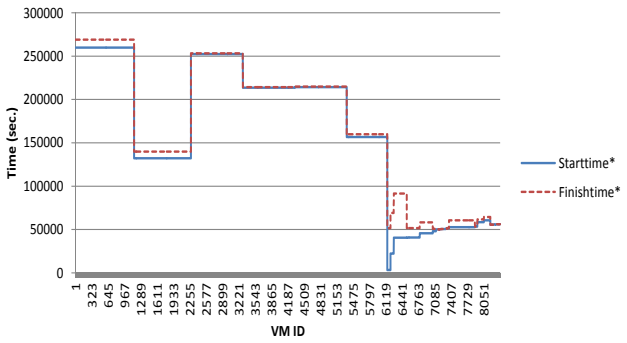

Fig. 3. Starting time (blue line) and finishing time (dotted red line) of VMs in simulations with SDSC Blue Horizon log-trace [13]. Both starting time* and finishing time* in the chart which has their values from the simulated starting and finishing time subtracts 390,000 seconds.

TABLE III
HOST POWER CONSUMPTION MODEL WITH CPU UTILIZATION OF AN TYPICAL SERVER WITH 4 CORES, 8192 MBYTES OF RAM, 10GB/S OF NETWORK BANDWIDTH, 1TBYTES OF STORAGE.

CPU Utilization (%)	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
Host power (Watts)	93.7	97.0	101.0	105.0	110.0	116.0	121.0	125.0	129.0	133.0	135.0

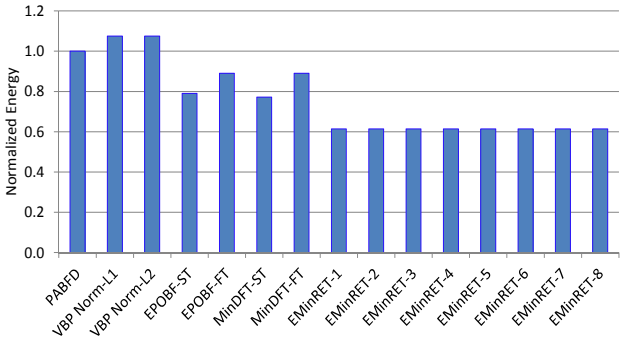

Fig. 4. Normalized energy. Result of simulations with SDSC Blue Horizon log-trace.

micro VM. Figure 1 shows chart of starting times and finishing times of the VMs in a simulation (the simulations have the same starting times and duration times of VMs). All physical machines are identical and each physical machine has x4 CPU cores (2660 MIPS per core), 8192 MB of physical memory, 10 Gb/s of network bandwidth, 1 TBytes of storage. Power model of each physical machine is shown in Table III, which is from a typical Hewlett-Packard Company ProLiant ML110 G5 in SPECpower_ssj2008 benchmark [26]. In the simulations, we use weights as following: (i) weight of increasing time of mapping a VM to a host: {0.001, 0.01, 1, 100, 3600}; (ii) weights of computing resources such as number of MIPS per CPU core, physical memory (RAM), network bandwidth, and storage respectively: 940, 24414, 1, 0.0001 respectively. We will discuss how to choose these values for weights of resources in another paper. We simulate on combination of these weights. The total energy consumption of each EMinRET-[1-8] is the average of five times simulation with various weights of increasing time (e.g. 0.001, 0.01, 1, 100, or 3600).

We choose PABFD [1] as the baseline algorithm because the PABFD is a famous power-aware best-fit decreasing in the energy-aware scheduling research community. We also compare our proposed VM allocation algorithms with two vector bin-packing algorithms (VBP-Norm-L1/L2) to show the importance of with/without considering VM's starting time and finish time in reducing the total energy consumption of VM placement problem.

C. Results and Discussions

The Table IV shows simulation results of scheduling algorithms solving scheduling problems with 5,687 VMs and 10,000 physical machines (hosts), in which VM's data is converted from the HPC2N Seth log-trace [14]. The Table V shows simulation results of scheduling algorithms solving scheduling problems with 8,368 VMs and 5,000 physical machines (hosts), in which VM's data is converted from the SDSC BLUE log-trace [13]. Both Figure 2 and Figure 4 show a bar chart comparing energy consumption of VM allocation algorithms that scale with the PABFD. None of the algorithms use VM migration techniques, and all of them satisfy the Quality of Service (e.g. the scheduling algorithm provisions maximum of user VM's requested resources). We use total energy consumption as the performance metric for evaluating these VM allocation algorithms. The energy saving shown in both Table IV and Table V are the reduction of total energy consumption of the corresponding algorithm compared with the baseline PABFD [1] algorithm.

Table IV shows that, compared with PABFD [1], our EMinRET with 8 configurations can reduce the total energy consumption by average 34.25%, and average 34.25% respectively compared with norm-based vector bin-packing algorithms (VBP-Norm-L1/L2) in simulations with the first 300 jobs of the HPC2N Seth log-trace. Table V shows that, compared with PABFD [1] or norm-based vector bin-packing algorithms (e.g. VBP-Norm-L2) [2], our EMinRET-[1-8] reduces the total energy consumption by average 39% in simulations with the first 50 jobs in the SDSC BLUE log-trace.

The PABFD generates a schedule that uses higher energy consumption than the MinRET-ST and EMinRET-FT because of the following main reasons. First, our hypothesis in this paper is that each VM consumes the same amount of energy in any physical server (e_i) and all physical servers are identical. In consequence, the PABFD will choose a random physical server to map a new VM. The PABFD sorts the list of VMs by decreasing requested computing power (e.g. MIPS), therefore the PABFD allocates VMs that has the most requested computing power firstly. In Table II, a VM1-typed VM has highest requested computing power in the list, next is VM2-typed VMs, etc.. Instead, our proposed EMinRET-[1-8] algorithms assign a new VM to a physical server in such a way that has minimum increase of completion times and use fully all resources in physical machines.

These EMinRET-[1-8] algorithms perform better than our previous algorithms such as MinDFT-ST/FT and EPOBF-ST/FT in the simulations. Compared to EPOBF-ST and EPOBF-FT, the EMinRET-[1-8] have less total energy consumption than by average 14.2% and respectively average 14.9%. The EMinRET-[1-8] have also less total energy consumption than the MinDFT-ST and MinDFT-FT from average 10.6% and respectively 14.9%. In the simulations, swapping between a new VM and its overlapped VM that is allocated to a host reduce total completion time on the host. For input as in Table I, the VM2 is remove from the first host, the VM6 will allocated to the first host.

VI. Conclusions and Future Work

In this paper, we formulated an energy-aware VM allocation problem with fixed starting time and non-preemption. We also discussed our two key observations in the VM allocation problem. First, minimizing total energy consumption is equivalent to minimizing the sum of total completion time of all physical machines (PMs). For some possible schedules, which have same total completion time of all PMs, the RET metric decides a schedule that has higher resource efficiency. Second, swapping between an unallocated VM and its overlapped VM, which has already been allocated to a PM, can reduce the total completion time of all PMs. Based on these observations, we proposed EMinRET algorithm to solve the energy-aware VM allocation with fixed starting time and duration time. Our proposed EMinRET and its sorting list of VMs by starting time (or longest duration time first, or latest finishing time first) can all reduce the total energy consumption of the physical servers compared with other algorithms in simulation results on the HPC2N Seth [14] and SDSC Blue Horizon log-traces [13] in the Feitelson's PWA [25]. The combination of EMinRET with its sorting list of VMs by latest finishing time first (EMinRET-8) is the best.

In future, we are developing EMinRET into a cloud resource management software (e.g. OpenStack Nova Scheduler). Additionally, we are working on IaaS cloud systems with heterogeneous physical servers and job requests consisting of multiple VMs. We are studying how to choose the right weights of time and resources (e.g. computing power, physical memory, network bandwidth, etc.) in Machine Learning techniques.

Acknowledgment

This research was conducted within the "Studying and developing practical heuristics for energy-aware virtual machine-based lease scheduling problems in cloud virtualized data centers" sponsored by TIS.

References

[1] A. Beloglazov and R. Buyya, "Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in cloud data centers," *Concurrency and Computation: Practice and Experience*, vol. 24, no. 13, pp. 1397–1420, 2012.

[2] R. Panigrahy, K. Talwar, L. Uyeda, and U. Wieder, "Heuristics for Vector Bin Packing," Microsoft Research, Tech. Rep., 2011.

[3] B. Sotomayor, "Provisioning Computational Resources Using Virtual Machines and Leases," Ph.D. dissertation, University of Chicago, 2010.

[4] A. Beloglazov, J. Abawajy, and R. Buyya, "Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing," *Future Generation Comp. Syst.*, vol. 28, no. 5, pp. 755–768, 2012.

[5] L. A. Barroso, J. Clidaras, and U. Hölzle, "The datacenter as a computer: An introduction to the design of warehouse-scale machines," *Synthesis lectures on computer architecture*, vol. 8, no. 3, pp. 1–154, 2013.

[6] X. Fan, W.-D. Weber, and L. Barroso, "Power provisioning for a warehouse-sized computer," in *ISCA*, 2007, pp. 13–23.

[7] K. Le, R. Bianchini, J. Zhang, Y. Jaluria, J. Meng, and T. D. Nguyen, "Reducing electricity cost through virtual machine placement in high performance computing clouds," in *SC*, 2011, p. 22.

[8] S. K. Garg, C. S. Yeo, A. Anandasivam, and R. Buyya, "Energy-Efficient Scheduling of HPC Applications in Cloud Computing Environments," *CoRR*, vol. abs/0909.1146, 2009.

[9] N. Quang-Hung, N. Thoai, and N. T. Son, "EPOBF: Energy Efficient Allocation of Virtual Machines in High Performance Computing Cloud," *TLDKS XVI*, vol. LNCS 8960, pp. 71–86, 2014.

[10] I. Takouna, W. Dawoud, and C. Meinel, "Energy Efficient Scheduling of HPC-jobs on Virtualize Clusters using Host and VM Dynamic Configuration," *Operating Systems Review*, vol. 46, no. 2, pp. 19–27, 2012.

[11] H. Viswanathan, E. K. Lee, I. Rodero, D. Pompili, M. Parashar, and M. Gamell, "Energy-Aware Application-Centric VM Allocation for HPC Workloads," in *IPDPS Workshops*, 2011, pp. 890–897.

[12] N. Quang-Hung, D.-K. Le, N. Thoai, and N. T. Son, "Heuristics for Energy-Aware VM Allocation in HPC Clouds," *Future Data and Security Engineering (FDSE 2014)*, vol. 8860, pp. 248–261, 2014.

[13] "The SDSC BLUE log-trace (SDSC-BLUE-2000-4.1-cln.swf.gz file)," (retrieved on 01 Nov. 2014), http://www.cs.huji.ac.il/labs/parallel/workload/l_sdsc_blue/SDSC-BLUE-2000-4.1-cln.swf.gz.

[14] "The HPC2N Seth log-trace (HPC2N-2002-2.2-cln.swf.gz file)," (retrieved on 01 May 2015), http://www.cs.huji.ac.il/labs/parallel/workload/l_hpc2n/HPC2N-2002-2.2-cln.swf.gz.

[15] T. Knauth and C. Fetzer, "Energy-aware scheduling for infrastructure clouds," in *4th IEEE International Conference on Cloud Computing Technology and Science Proceedings*. IEEE, Dec. 2012, pp. 58–65.

[16] L. Chen and H. Shen, "Consolidating complementary VMs with spatial/temporal-awareness in cloud datacenters," in *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*. IEEE, Apr. 2014, pp. 1033–1041.

[17] A. Beloglazov, R. Buyya, Y. C. Lee, and A. Zomaya, "A Taxonomy and Survey of Energy-Efficient Data Centers and Cloud Computing Systems," *Advances in Computers*, vol. 82, pp. 1–51, 2011.

[18] A.-C. Orgerie, M. D. de Assuncao, and L. Lefevre, "A survey on techniques for improving the energy efficiency of large-scale distributed systems," *ACM Computing Surveys*, vol. 46, no. 4, pp. 1–31, Mar. 2014.

[19] A. Hameed, A. Khoshkbarforoushha, R. Ranjan, P. P. Jayaraman, J. Kolodziej, P. Balaji, S. Zeadally, Q. M. Malluhi, N. Tziritas, A. Vishnu, S. U. Khan, and A. Zomaya, "A survey and taxonomy on energy efficient resource allocation techniques for cloud computing systems," pp. 1–24, Jun. 2014.

[20] T. Mastelic, A. Oleksiak, H. Claussen, I. Brandic, J.-M. Pierson, and A. V. Vasilakos, "Cloud computing: Survey on energy efficiency," *ACM Comput. Surv.*, vol. 47, no. 2, pp. 33:1–33:36, Dec. 2014.

[21] M. Y. Kovalyov, C. Ng, and T. E. Cheng, "Fixed interval scheduling: Models, applications, computational complexity and algorithms," *European Journal of Operational Research*, vol. 178, no. 2, pp. 331–342, 2007.

[22] E. Angelelli and C. Filippi, "On the complexity of interval scheduling with a resource constraint," *Theoretical Computer Science*, vol. 412, no. 29, pp. 3650–3657, 2011.

[23] M. Flammini, G. Monaco, L. Moscardelli, H. Shachnai, M. Shalom, T. Tamir, and S. Zaks, "Minimizing total busy time in parallel scheduling with application to optical networks," *Theoretical Computer Science*, vol. 411, no. 40-42, pp. 3553–3562, Sep. 2010.

[24] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. F. De Rose, and R. Buyya, "Cloudsim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms," *Softw., Pract. Exper.*, vol. 41, no. 1, pp. 23–50, 2011.

[25] "Feitelson's Parallel Workloads Archive," (retrieved on 31 Januray 2014), http://www.cs.huji.ac.il/labs/parallel/workload/.

[26] "Hewlett-Packard ProLiant ML110 G5 Power on SPECpowers-sj2008 Benchmark," (retrieved on 30 Jan 2015), http://www.spec.org/power_ssj2008/results/res2011q1/power_ssj2008-20110124-00339.html.